

## ***NASA Center for Climate Simulation Success Story***

### **NCCS Triples Supercomputer Performance for Earth Science Modeling**

April 28, 2015

The NASA Center for Climate Simulation (NCCS) is nearly tripling the peak performance of its Discover supercomputer to more than 3.3 petaflops, or 3,361 trillion floating-point operations per second. This unprecedented NCCS upgrade is necessary to meet the exploding demands of NASA's Earth science modeling efforts.

A rigorous open procurement included the benchmarking of software codes run at NCCS, especially the Goddard Earth Observing System Model, Version 5 (GEOS-5) and the NASA Unified-Weather and Research Forecasting (NU-WRF) Model. The solution offering the best value was SGI® Rackable® clusters, which are replacing portions of Discover dating from 2011. NCCS is installing the SGI Rackable hardware as three Scalable Compute Units (SCUs) using current-generation 14-core Intel E5-2697v3 (Haswell) processors. Together the new SCUs will offer a total of 64,512 processor cores.

Supercomputer installations require choreography between NCCS technical and facilities staff and the computer vendor. The dance becomes more intricate when replacing existing SCUs. "We want to have the old hardware out at least a week beforehand," said Bruce Pfaff, who leads Discover's system administration team. "But we also want to maximize the amount of time users have with the old system and minimize the period of limited resources during the installation."

The prelude involves planning for nearly 1 megawatt of power and 400 tons of cooling, ensuring the vendor factory configures the racks for optimal onsite operations, and acquiring 10 nodes for the NCCS Test and Development System (TDS) to prepare the operating system (OS) and software stack (compilers, file system, MPI, etc.). System administrators also scrub all NASA data off the old hardware before physically removing it.

Installation Act 1 stars the vendor. SGI staff positioned the racks, connected the power supplies and chilled-water cooling pipes, threaded hundreds of cables between racks, and ran benchmarks to check the hardware. With SGI Rackable being a new architecture for NCCS, center staff played substantial supporting roles with the first and biggest unit—SCU10. Notably they and SGI spent about a week balancing power for the 30,240-core cluster.

In Act 2 NCCS system administrators take center stage. Their main tasks are building interconnections with the rest of the supercomputer and installing the OS and software stack. TDS preparations were instrumental since SCU10 required a new OS. System administrators also developed software to address issues with the node management system and upgraded the InfiniBand network firmware.

The Act 3 spotlight swings to the NCCS benchmarking team and NASA pioneer users to shake out any other hardware and software issues. SCU10 hosted an ultra-high-resolution GEOS-5 simulation that consumed the entire cluster (this simulation will be the subject of a future NCCS Success Story). "Once our power users put the system through its paces, we are pretty confident with its stability and robustness for the general user population," Pfaff said.

NCCS kept its user community's needs constantly in mind when tuning the supercomputer architecture with SGI. "As model codes are scaling to higher and higher core counts, they need larger amounts of memory," said NCCS High-Performance Computing Lead Daniel Duffy. Extending a practice begun with SCU9, all three SCUs have over 4 gigabytes of memory per core.



*The newest units of the NASA Center for Climate Simulation (NCCS) Discover supercomputer are SGI Rackable clusters that will house a total of 64,512 processor cores. Photo by Bruce Pfaff.*

SCU10's 138 terabytes of memory are proving handy for its primary user—the Downscaling Project. The NASA-wide investigation is assessing the credibility of downscaled climate projections, which use the results from global climate models to drive higher-resolution regional models. For this project the global model is GEOS-5 running at 12 kilometers (km), and the regional model is NU-WRF running at 24, 12, and 4 km. Scientists are comparing how well the models predict three weather phenomena impacting the continental United States: Northeast wintertime storms, midcontinent summertime storms, and West Coast wintertime atmospheric rivers.

Another computer architecture advantage for NASA modeling efforts is a fully non-blocking interconnect fabric. On the new SCUs each 28-core node can communicate directly with every other node via Fourteen Data Rate (FDR) InfiniBand rated at 56 gigabits per second. “The research runs for GEOS-5 are getting very large and require extremely fast computations and communications. This is a major challenge if the high-performance computing fabric is not suited to those runs,” Duffy said.

Since a high-resolution simulation can produce up to several petabytes of data, NCCS is more than doubling Discover's online disk to 33 petabytes.

SCU10 became available to general users in January, and SCU11 is currently in pioneer user mode. NCCS expects SCU12 to arrive in late May. A normal NCCS installation pace is one SCU per year. “By the time we're done, we will have installed three SCUs in 7 months,” said Discover system administrator Mike Donovan.

Other Discover system administrators are Nicko Acks, Jim Carlisi, Michael Chyatte, Lyn Gerner, Dave Kemeza, Aaron Knister, Jonathan Mills, and Jordan Robertson. Facilities experts are Steve Majstorovic and Hal Domchick. Benchmarking team members are Duffy and Eric Winter.

*Jarrett Cohen, NASA Goddard Space Flight Center*

### Contacts

Dan Duffy  
High-Performance Computing Lead  
NASA Center for Climate Simulation  
NASA Goddard Space Flight Center  
daniel.q.duffy@nasa.gov  
301.286.8830


Bruce Pfaff  
Lead System Administrator, Discover Supercomputer  
NASA Center for Climate Simulation  
NASA Goddard Space Flight Center  
bruce.e.pfaff@nasa.gov  
301.286.8587

### More Information

Discover Supercomputer  
[http://www.nccs.nasa.gov/discover\\_front.html](http://www.nccs.nasa.gov/discover_front.html)



*As with previous supercomputer upgrades, NCCS staff worked closely with vendor SGI on installing the new racks. Photos by Bruce Pfaff and Michael Chyatte.*



**NASA Center for Climate Simulation**  
*High-Performance Science*  
<http://www.nccs.nasa.gov>